

Signaling in Quasipolynomial Time

Yu Cheng[‡]

Ho Yee Cheung^{*}

Shaddin Dughmi[†]

Shang-Hua Teng[‡]

Department of Computer Science
University of Southern California

Abstract

Strategic interactions often take place in an environment rife with uncertainty. As a result, the equilibrium of a game is intimately related to the information available to its players. The *signaling problem* abstracts the task faced by an informed “market maker”, who must choose how to reveal information in order to effect a desirable equilibrium.

In this paper, we consider two fundamental signaling problems: one for abstract normal form games, and the other for single item auctions. For the former, we consider an abstract class of objective functions which includes the social welfare and weighted combinations of players’ utilities, and for the latter we restrict our attention to the social welfare objective and to signaling schemes which are constrained in the number of signals used. For both problems, we design approximation algorithms for the signaling problem which run in quasi-polynomial time under various conditions, extending and complementing the results of various recent works on the topic.

Underlying each of our results is a “meshing scheme” which effectively overcomes the “curse of dimensionality” and discretizes the space of “essentially different” posterior beliefs – in the sense of inducing “essentially different” equilibria. This is combined with an algorithm for optimally assembling a signaling scheme as a convex combination of such beliefs. For the normal form game setting, the meshing scheme leads to a convex partition of the space of posterior beliefs and this assembly procedure is reduced to a linear program, and in the auction setting the assembly procedure is reduced to submodular function maximization.

^{*}Supported in part by USC Provost’s Ph.D. Fellowship and NSF CCF-0964481.

[†]Supported in part by NSF CAREER Award CCF-1350900.

[‡]Supported in part by NSF CCF-0964481 and NSF CCF-1111270.

1 Introduction

In games with uncertainty, the information available to players may influence their strategic decisions, and could fundamentally impact their utilities and their beliefs regarding the preferences of other players. Consequently, a principal (market maker) who is privy to confidential state-of-nature information may affect the equilibrium by strategically “leaking” information ahead of players’ decisions — a practice often referred to as *signaling*. If the principal’s goal is to boost her own objective in the outcome of the game, then her strategic decision regarding which information to reveal is an optimization problem, which we refer to as the *signaling problem*.

We consider *Bayesian games* in which players have *prior* beliefs on the structure of the game, and form *posterior beliefs* based on the information revealed by the principal. We restrict our attention to games in which players share the same prior belief, and *symmetric signaling schemes* which reveal the same information to all players. Thus, a signaling scheme defines a number of new Bayesian (sub-)games — one for each realization of the information string (aka the *signal*). From this viewpoint, the search space for signaling is the space of decompositions of the prior belief into posterior beliefs. The goal of the principal is then to identify a signaling scheme which induces, on average over the resulting Bayesian subgames, a favorable equilibrium.

We focus on two fundamental models previously considered in a signaling context. Our first model, considered in [Dug14], is the most abstract: normal form games in which the payoff matrix is parametrized by a state of nature, which is drawn from a common prior. We consider the signaling question in normal form games for an abstract class of objectives which include the social welfare, weighted combinations of players’ utilities, and utility functions which depend on players’ actions in the various states of nature. Our second model, considered in [EFG⁺12, BMS12, DIR14] is that of a Bayesian second-price auction, in particular when there is uncertainty regarding the identity or attributes of the item for sale, and players’ private valuations for the various configurations of the item are drawn from a distribution. We restrict attention to the social welfare objective in this model, yet as in [DIR14] we place a bound on the communication bandwidth of the signaling scheme — equivalently the number of different signals used to describe the item.

1.1 Context and Results

There are significant computational challenges involved in obtaining an (approximately) optimal signaling scheme. Whereas the space of the posterior beliefs is a convex set, signaling problems in even very simple settings involve a non-convex objective, or combinatorial constraints, which either lead to hardness results or require novel algorithmic approaches. This is borne out in a number of recent works which examine the algorithmic aspects of signaling, and we mention those most relevant to this paper next.

Signaling in the Bayesian second-price auction model has perhaps received the most attention from an algorithmic perspective. Emek et al. [EFG⁺12] and Miltersen and Sheffet [BMS12], consider the signaling problem faced by an informed auctioneer who must decide to (partially) reveal information regarding the item for sale, with the goal of maximizing his revenue. Whereas the general problem is shown to be NP-hard, the special case in which players’ valuation distribution has small support is solved in polynomial time. We also mention the work of Guo and Deligkas [GD13], who consider a similar question albeit in a different model of uncertainty, and similarly present NP-hardness results and heuristics for their model.

More recently, Dughmi et al. [DIR14] consider the same auction setting, albeit with the welfare objective and subject to exogenous constraints on the signaling policy — most notably a communication constraint. There the problem was shown to be NP-hard to approximate better than a

$(1-1/e)$ factor, and approximation algorithms matching this guarantee were presented in the special case in which players’ valuations are drawn from a small support distribution. The present paper considers the general algorithmic question without a bound on players’ valuation distribution, while relaxing the runtime requirement to quasi-polynomial. Specifically, we derive a quasi-polynomial-time algorithm with an approximation ratio of $(1 - 1/e)$, modulo an additive loss of ϵ times the upper limit of a player’s valuation. Our result holds when players’ valuation distribution is given explicitly, and more generally whenever said distribution can be sampled efficiently.

The other signaling model considered in this paper, that of abstract normal form games, has only recently been examined from a complexity-theoretic perspective. Dughmi [Dug14] considered the special case of two-player zero-sum games, and examined the design of symmetric signaling schemes with the goal of maximizing the expected utility of one of the players. The results were mostly negative, assuming the conjectured hardness of the planted clique problem. Specifically, it was shown that no fully polynomial time approximation scheme (FPTAS) is possible for the signaling problem for zero sum games, assuming the planted clique conjecture. This leaves open the potential for a PTAS. In the present paper, we complement the impossibility result of [Dug14] with a quasi-polynomial time approximation scheme (QPTAS), which applies more generally to general sum normal form games with a constant number of players, and an abstract class of objectives which includes the social welfare and weighted combinations of players’ utilities as a special case. For general sum games, our result is a bicriteria result which forgoes the exact Nash equilibrium, and instead computes a signaling scheme with ϵ -equilibria that are competitive with the ϵ' -equilibria of every other signaling scheme whenever $\epsilon' < \epsilon$ (including $\epsilon' = 0$). Our result can also be extended to Stackelberg games.

1.2 Our Techniques

Both our results rely on a “meshing method” applied to the space of posterior beliefs, or equivalently the equilibria which they induce. The method proceeds in two main steps.

1. *Construct a Small Dictionary of Equilibria:* This is a discrete family of objects which indexes the potential equilibria of a signaling scheme, with the property that they form an “ ϵ -net” of the space of all equilibria in a precise technical sense. This “net” forms our signaling “dictionary”.
2. *Construct a near-optimal Signaling Scheme:* We then assemble a near-optimal signaling scheme which induces subgames with equilibria from our dictionary. This involves solving a nontrivial optimization problem which optimally decomposes the prior distribution into posterior beliefs inducing equilibria in our dictionary.

The technical challenge involved in step (1) is the identification of a net of the space of equilibria which has “low approximation dimensionality” with respect to the space of signaling schemes and the design objective. This arises due to the “curse of dimensionality” of the space of posterior beliefs and the space of equilibria. For step (2), the challenge arises from the fact that not every convex combination of posterior beliefs is a valid signaling scheme — indeed, the signaling scheme must be a convex decomposition of the prior distribution. This induces a nontrivially constrained optimization problem.

For the results of Section 3 for normal form games, our dictionary is based on a Lemma of Lipton et al. [LMM03]. Specifically, their result implies the existence of a quasi-polynomial-sized family of mixed strategy profiles which, simultaneously for all games and equilibria of those games, includes a profile which approximates the payoffs of the equilibrium within an additive ϵ , and itself

forms an ϵ -equilibrium. To combine these approximate equilibria into a signaling scheme, we make two observations: First, the space of posterior beliefs inducing a particular equilibrium forms a convex polytope; Second, the optimization problem of optimally partitioning the prior belief into a quasi-polynomial number of posterior beliefs, one in each polytope corresponding to an equilibrium, can be formulated via a linear program after an appropriate change of variables.

For the results of Section 4 for second-price auctions, constructing our dictionary takes more work. Specifically, we consider all rules for selecting a winner of the auction as a function of the drawn valuation matrix, and show that a quasi-polynomial number of such rules suffices to approximate the welfare in all equilibria to within an additive error of ϵ . Assembling a signaling scheme which induces such winner selection rules in an optimal combination is then reduced to submodular function maximization subject to a cardinality constraint, which admits a $(1 - 1/e)$ -approximation algorithm.

1.3 Additional Discussion of Related Work

The study of the effects of information on strategic interactions, and mechanisms for signaling, has its roots in the early works of Akerlof [Ake70] and Spence [Spe73]. Hirshleifer [Hir71] was the first to observe that optimal information revelation is nontrivial, in the sense that more information sometimes leads to worse market outcomes. This contrasts with earlier work by Blackwell [Bla51] which implied that more information is always better for a single agent in a non-competitive environment. Since then, many works have examined the effects of additional information on players' equilibrium utilities. Lehrer et al. [LRS10] showed that additional information improves players' utilities in *common interest* games — games where players have identical payoffs in each outcome, and Bassan et al. [BGSZ03] exhibited a polyhedral characterization of games in which more information improves the individual utility of *every* player. The “linkage principle” of Milgrom and Weber [MW82] exhibits natural conditions under which revealing additional information to buyers in an auction improves the seller's revenue; and Syrgkanis et al. [SKT13] examine failings of the linkage principle in common value auctions. Unlike revenue, it is known that additional information always improves welfare in a second-price auction.

Despite appreciation of the importance of information in strategic interactions, it is only recently that researchers have started viewing the information structure of a game as a mathematical object to be designed, rather than merely an exogenous variable. Kamenica and Gentzkow [KG09] examine settings in which a sender must design a signaling scheme to convince a less informed receiver to take a desired action. Recent work in the CS community, including by Emek et al. [EFG⁺12], Miltersen and Sheffet [BMS12], and Guo et al. [GD13], examines revenue-optimal signaling in an auction setting, and presents polynomial-time algorithms and hardness results for computing it. Dughmi et al. [DIR14] examine welfare-optimal signaling in an auction setting under exogenous constraints, and presents polynomial-time algorithms and hardness results under assumptions on players' value distributions; our auction signaling result applies to a generalization of their model. Also relevant to our work is the recent result of Dughmi [Dug14], which essentially rules out an FPTAS for the signaling problem in zero-sum games; our result on general games is a positive counter-point to the result of [Dug14], slightly relaxing both the equilibrium definition and the polynomial-time restriction.

2 Preliminaries: Information Revelation and Signaling Schemes

2.1 Games

A *Bayesian game* is a family of *games of complete formation* parametrized by a state of nature θ , where θ is assumed to be drawn from a known prior distribution. We consider two classes of such games: *explicit normal form Bayesian games*, and *Bayesian single item auction games*.

Normal Form games

We consider normal form games of incomplete information, given by the following parameters.

- A positive integer n denoting the number of players, which we index by the set $[n] = \{1, \dots, n\}$.
- A nonnegative integer m bounding the number of pure strategies of each player.¹
- A finite family $\Theta = \{1, \dots, M\}$ of *states of nature*, which we index by θ .
- A family of *payoff tensors* $\mathcal{A}_i^\theta : [m]^n \rightarrow [-1, 1]$, one per player i and state of nature θ , where $\mathcal{A}_i^\theta(s_1, \dots, s_n)$ is the payoff to player i when the state of nature is θ and each player j plays strategy s_j .

A *Bayesian* normal form game of incomplete information is additionally equipped with a *common prior distribution* $\lambda \in \Delta_M$ on the states of nature. Absent additional informational, risk neutral players behave as in the complete information game $\mathbf{E}_{\theta \sim \lambda}[\mathcal{A}^\theta]$. We consider signaling schemes which partially and symmetrically inform players by publicly announcing a signal σ , correlated with θ ; this induces a common posterior belief on the state of nature for each value of σ . When players' posterior belief over θ is given by $\mu \in \Delta_M$, we use \mathcal{A}^μ to denote the equivalent complete information game $\mathbf{E}_{\theta \sim \mu}[\mathcal{A}^\theta]$. As shorthand, we use $\mathcal{A}_i^\mu(x_1, \dots, x_n)$ to denote $\mathbf{E}[\mathcal{A}_i^\theta(s_1, \dots, s_n)]$ when $\theta \sim \mu \in \Delta_M$ and $s_i \sim x_i \in \Delta_m$. In the event that the state of nature is θ and players play the pure strategy profile s_1, \dots, s_n , we refer to the tuple $(\theta, s_1, \dots, s_n)$ as the *state of play*.

For all our results, we assume that a Bayesian game (\mathcal{A}, λ) is represented explicitly as a list of tensors $\{\mathcal{A}_i^\theta \in [-1, 1]^{m^n} : i \in [n], \theta \in [M]\}$, and a vector $\lambda \in \Delta_M$.

Single Item Second-Price Auctions

We consider Bayesian second-price auctions described by the following parameters:

- A nonnegative integer n denoting the number of bidders. We index the bidders by the set $[n] = \{1, \dots, n\}$.
- A finite family $\Theta = \{1, \dots, M\}$ of *states of nature*, indexed by θ . The states of nature represent potential configurations of the item for sale.
- A common-knowledge prior distribution $\lambda \in \Delta_M$ on the states of nature.
- A common-knowledge prior distribution \mathcal{D} on *valuation matrices* $\mathcal{V} \in [0, 1]^{n \times M}$

¹Without loss of generality, we assume each player has exactly m pure strategies.

We assume the state of nature $\theta \in \Theta$, describing the item being sold in the auction, is first drawn from λ , and revealed to the auctioneer but not to the bidders. Then the auctioneer reveals a public signal σ , a (partial) description of the item, according to some signaling scheme (See Section 2.2). Subsequently, players' valuations $\mathcal{V} \in [0, 1]^{n \times M}$ are drawn from \mathcal{D} , where $\mathcal{V}_{i\theta}$ is player i 's value for the item described by θ , and each player i privately learns his valuation \mathcal{V}_i . Finally a second-price auction is run, where bidders bid according to their private valuations and their posterior belief (after learning σ) regarding the configuration of the item being sold.²

We note that we assume that \mathcal{D} and λ are independent. We also emphasize that the auctioneer knows nothing regarding \mathcal{V} besides its distribution \mathcal{D} prior to running the auction, and that the bidders know nothing regarding θ besides its distribution λ and the signal σ .

For our algorithmic results, we assume that λ is represented explicitly as a vector, and that \mathcal{D} is given explicitly as a list of $n \times M$ valuation matrices $\mathcal{V}^1, \dots, \mathcal{V}^r \in [0, 1]^{n \times M}$ with associated probabilities ρ_1, \dots, ρ_r summing to 1. However, as we show in Section 4.2, our results also generalize to the case in which \mathcal{D} can only be sampled efficiently. The latter includes, as a special case, the scenario in which players' valuations are independent of each other, with each represented explicitly as a list of vectors in $[0, 1]^M$ with associated probabilities.

2.2 Signaling Schemes

We examine policies whereby a principal reveals partial information regarding the state of nature θ to the players. For our main results, we require that the principal reveal the same information to all players in the game. A *symmetric signaling scheme* is given by a set Σ of *signals*, and a (possibly randomized) map φ from states of nature Θ to signals Σ . Abusing notation, we use $\varphi(\theta, \sigma)$ to denote the probability of announcing signal $\sigma \in \Sigma$ conditioned on the state of nature is $\theta \in \Theta$. We restrict attention to signaling schemes with a finite set of signals Σ , and this is without loss of generality when Θ is finite. We elaborate on this after describing the convex decomposition interpretation of a signaling scheme.

We note that signaling schemes are in one-to-one correspondence with *convex decompositions* of the prior distribution $\lambda \in \Delta_M$ — namely, distributions supported on the simplex Δ_M , and having expectation λ . Formally, a signaling scheme $\varphi : \Theta \rightarrow \Sigma$ corresponds to the convex decomposition $\lambda = \sum_{\sigma \in \Sigma} \alpha_\sigma \cdot \mu_\sigma$, where $\alpha_\sigma = \mathbf{Pr}[\varphi(\theta) = \sigma] = \sum_{\theta \in \Theta} \lambda(\theta) \varphi(\theta, \sigma)$, and $\mu_\sigma(\theta) = \mathbf{Pr}[\theta | \varphi(\theta) = \sigma] = \frac{\lambda(\theta) \varphi(\theta, \sigma)}{\alpha_\sigma}$. Note that α_σ is the probability of announcing signal σ , and $\mu_\sigma \in \Delta_M$ is the *posterior distribution* of θ conditioned on signal σ . The converse is also true: every convex decomposition of λ corresponds to a signaling scheme.

We judge the quality of a signaling scheme by the outcome it induces signal by signal. Specifically, the principal is equipped with an objective function of the form $\sum_{\sigma} \alpha_\sigma \cdot f(\mu_\sigma)$, where $f : \Delta_M \rightarrow \mathbb{R}$ is some function mapping a posterior distribution to the quality of the equilibrium chosen by the players. For example, f may be the social welfare at the induced equilibrium, or any weighted combination of players' utilities at equilibrium, or something else entirely. In this setup, one can show that there always exists an signaling scheme with a finite set of signals which maximizes our objective, so long as the states of nature are finitely many. The optimal choice of signaling scheme is related to the *concave envelope* f^+ of the function f .³ Specifically, such a signaling scheme achieves $\sum_{\sigma} \alpha_\sigma \cdot f(\mu_\sigma) = f^+(\lambda)$. Application of Caratheodory's theorem to the hypograph of f , therefore, shows that $M + 1$ signals suffice.

²As mentioned in [DIR14], the particular choice of auction is immaterial.

³ f^+ is the point-wise lowest concave function h for which $h(x) \geq f(x)$ for all x in the domain. Equivalently, the hypograph of f^+ is the convex hull of the hypograph of f .

For concreteness, the reader can think of a signaling scheme φ as represented by the matrix of pairwise probabilities $\varphi(\theta, \sigma)$. Since we only consider games where the states of nature, and therefore also the number of signals w.l.o.g., are polynomially many in the description size of the game, this is a compact representation. The representation of φ as a convex decomposition would do equally well, as both representations can be efficiently computed from each other.

2.3 Meshing Schemes at a High Level

Note that direct discretization of the space of posterior beliefs Δ_M is usually impractical when M is large. Instead, our meshing scheme overcomes this “curse of dimensionality” by using a partition of Δ_M into a quasi-polynomial number of subspaces (of posterior beliefs). We then select at most one posterior distribution from each component of the partition as the convex decomposition of the prior λ . Naturally, such partition of Δ_M needs to satisfy both the mathematical property that there exists an approximately optimal convex decomposition of λ that respects the partition, and the algorithmic property that such convex decomposition can be computed efficiently.

For Bayesian games, we use a net of the equilibrium space known to exist due to Lipton, Markakis, and Mehta [LMM03] to induce the desired partition of Δ_M , and use linear programming to compute the convex decomposition of λ with respect to the partition. For Bayesian auctions, we prove that logarithmic sized subsets of items can be used to define the desired partition of Δ_M that can be used in the submodular function maximization formulation developed in [DIR14].

2.4 Equilibria and Objectives

Given a Bayesian game, a symmetric signaling scheme (α, μ) with signals Σ induces $|\Sigma|$ subgames, one for each signal. The subgame corresponding to signal $\sigma \in \Sigma$ is played with probability α_σ , and players’ (common) beliefs regarding the state of nature in this subgame are given by the posterior distribution $\mu_\sigma \in \Delta_M$. The quality of a symmetric signaling scheme in such a game is contingent on a choice of an *equilibrium concept*, an *equilibrium selection rule*, and an *objective function*.

Normal Form Games

For normal form games, we adopt the approximate Nash equilibrium as our equilibrium concept. There are two variants.

Definition 2.1. *Let $\epsilon \geq 0$. In an n -player m -action normal form game with expected payoffs in $[-1, 1]$ given by tensors $\mathcal{A}_1, \dots, \mathcal{A}_n$, a mixed strategy profile $x_1, \dots, x_n \in \Delta_m$ is an ϵ -Nash Equilibrium (ϵ -NE) if*

$$\mathcal{A}_i(x_1, \dots, x_n) \geq \mathcal{A}_i(t_i, x_{-i}) - \epsilon$$

for every player i and alternative pure strategy $t_i \in [m]$.

Definition 2.2. *Let $\epsilon \geq 0$. In an n -player m -action normal form game with expected payoffs in $[-1, 1]$ given by tensors $\mathcal{A}_1, \dots, \mathcal{A}_n$, a mixed strategy profile $x_1, \dots, x_n \in \Delta_m$ is an ϵ -well-supported Nash equilibrium (ϵ -WSNE) if*

$$\mathcal{A}_i(s_i, x_{-i}) \geq \mathcal{A}_i(t_i, x_{-i}) - \epsilon$$

for every player i , strategy s_i in the support of x_i , and alternative pure strategy $t_i \in [m]$.

Clearly, every ϵ -WSNE is also an ϵ -NE. Note that we omitted reference to the state of nature in the above definitions — in a subgame corresponding to posterior beliefs $\mu \in \Delta_M$, we naturally use tensors $\mathcal{A}_1^\mu, \dots, \mathcal{A}_n^\mu$ instead.

Fixing an equilibrium concept (ϵ -NE or ϵ -WSNE), a Bayesian game (\mathcal{A}, λ) , and a signaling scheme $\varphi : \Theta \rightarrow \Delta(\Sigma)$, an *equilibrium selection rule* distinguishes an equilibrium strategy profile $(x_1^\sigma, \dots, x_n^\sigma)$ to be played in each subgame σ — we call the tuple $X = \{x_i^\sigma : \sigma \in \Sigma, i \in [n]\}$ a *Bayesian equilibrium* of the game (\mathcal{A}, λ) with signaling scheme φ . Together with the prior λ , the Bayesian equilibrium X induces a distribution $\Gamma \in \Delta_{\Theta \times [m]^n}$ over states of play — we refer to Γ as a *distribution of play*. We say Γ is *implemented* by signaling scheme φ in ϵ -NE (ϵ -WSNE). This is analogous to implementation of allocation rules in traditional mechanism design.

Our results concern objectives which depend only on the state of play, and we seek to maximize the objective in expectation over the distribution of play. These include, but are not restricted to, the social welfare of the players, as well as weighted combinations of player utilities. Formally, our objective is given by a tensor $\mathcal{F} : \Theta \times [m]^n \rightarrow [-1, 1]$.⁴ We seek a signaling scheme $\varphi : [M] \rightarrow \Sigma$, as well as a Bayesian ϵ -NE (ϵ -WSNE) $X = \{x_i^\sigma \in \Delta_m : i \in [n], \sigma \in \Sigma\}$, maximizing $\mathbf{E}[\mathcal{F}(\theta, \vec{s})]$ over the resulting distribution of play. We use $OPT_{NE}^\epsilon(\mathcal{A}, \lambda)$ and $OPT_{WSNE}^\epsilon(\mathcal{A}, \lambda)$ to denote the optimal values of these two optimization problems. As shorthand, we use $f(\varphi, X) = \mathbf{E}_{\theta \sim \lambda} \mathbf{E}_{\sigma \sim \varphi(\theta)} \mathbf{E}_{\vec{s} \sim x^\sigma}[\mathcal{F}(\theta, \vec{s})]$ to denote the expectation of our objective \mathcal{F} for signaling scheme φ and corresponding Bayesian equilibrium X .

Single Item Auctions

In our second-price auction game, we adopt the (unique) dominant-strategy truth-telling equilibrium as our solution concept. Specifically, given a symmetric signaling scheme $\varphi : \Theta \rightarrow \Sigma$, for the subgame corresponding to the signal $\sigma \in \Sigma$ it is a dominant strategy for player i to bid $\mathbf{E}_{\theta \sim \lambda}[\mathcal{V}_{i\theta} | \varphi(\theta) = \sigma]$ — his posterior expected value for the item conditioned on the received signal σ .

For our objective, we restrict attention to the *social welfare* — the expected value of the winning bidder for the item they win. Given a signaling scheme φ and a valuation matrix $\mathcal{V} \in [0, 1]^{n \times M}$, assuming bidders play the truth-telling equilibrium this is given by:

$$\begin{aligned} \text{welfare}(\varphi, \lambda, \mathcal{V}) &= \mathbf{E}_{\theta \sim \lambda} \mathbf{E}_{\sigma \sim \varphi(\theta)} \left[\max_{i=1}^n \mathbf{E}_{\theta \sim \lambda} [\mathcal{V}_{i\theta} | \varphi(\theta) = \sigma] \right] \\ &= \sum_{\sigma \in \Sigma} \max_{i=1}^n \sum_{\theta \in \Theta} \lambda(\theta) \varphi(\theta, \sigma) \mathcal{V}_{i\theta} \end{aligned}$$

When \mathcal{V} is drawn from a prior distribution \mathcal{D} , the expected welfare is given by

$$\text{welfare}(\varphi, \lambda, \mathcal{D}) = \mathbf{E}_{\mathcal{V} \sim \mathcal{D}} [\text{welfare}(\varphi, \lambda, \mathcal{V})].$$

As noted in [DIR14], convexity of the welfare in the probabilities $\varphi(\theta, \sigma)$ implies that full-information-revelation is optimal, in general. However, we consider signaling subject to a communication constraint — i.e. with $\Sigma = \{1, \dots, k\}$ for an input parameter $k < M$ limiting the number of different messages describing the item for sale.

3 Signaling in Normal Form Games

We now consider signaling in explicitly represented games when the adopted solution concept is the ϵ -Nash equilibrium or the ϵ -well-supported Nash equilibrium. We prove the following bi-criteria result.

⁴Equivalently, we may think of the objective as the payoff tensor of an additional player in the game.

Theorem 3.1. Fix $\epsilon > 0$, $\delta \geq 0$. Given as input an explicitly-described Bayesian normal form game (\mathcal{A}, λ) with $n = O(1)$ players, m actions, and M states of nature, and an objective $\mathcal{F} : [M] \times [m]^n \rightarrow [-1, 1]$, there is an algorithm with runtime $\text{poly}(M, m^{\ln m / \epsilon^2})$ which outputs a signaling scheme φ and corresponding Bayesian $(\epsilon + \delta)$ -equilibrium X satisfying $f(\varphi, X) \geq \text{OPT}^\delta(\mathcal{A}, \lambda) - \epsilon$. This holds for both approximate NE and approximate WSNE.

When the number of players is constant, we can in quasi-polynomial time approximate the optimal reward from signaling while losing an additive ϵ in the objective as well as in the incentive constraints. Our proof of this theorem hinges on three main lemmas: the first is drawn from the work of Lipton et al. [LMM03] regarding the existence of a quasi-polynomial-sized “net” of the space of equilibria; the second lemma states that the posterior beliefs implementing a particular approximate equilibrium form a simple polytope, in doing so reducing the signaling problem to optimization over convex decompositions of λ into a family of posteriors, each belonging to a given polytope; and the third lemma shows that optimization over such convex decompositions reduces to a linear program.

Lemma 3.2 (Lipton, Markakis and Mehta [LMM03]). Fix $n = O(1)$. For each integer m and $\epsilon > 0$, there is a family of mixed strategies $\mathcal{N}_{m, \epsilon} \subseteq \Delta_m$ with $|\mathcal{N}_{m, \epsilon}| \leq m^{O(\log m / \epsilon^2)}$ such that for every n -player m -action normal form game \mathcal{A} , objective tensor $\mathcal{F} : [m]^n \rightarrow [-1, 1]$, and δ -NE (WSNE) x of \mathcal{A} , there is an $(\epsilon + \delta)$ -NE (WSNE) y of \mathcal{A} such that each $y_i \in \mathcal{N}_{m, \epsilon}$, and $|\mathbf{E}_{s \sim x}[\mathcal{F}(s)] - \mathbf{E}_{s \sim y}[\mathcal{F}(s)]| \leq \epsilon$. Moreover, $\mathcal{N}_{m, \epsilon}$ can be enumerated in time $m^{O(\log m / \epsilon^2)}$.

Lemma 3.3. Fix a normal form game of incomplete information \mathcal{A} with n players, m actions, and M states of nature. Consider a mixed strategy profile $x = (x_1, \dots, x_n)$ with $x_i \in \Delta_m$. For each $\epsilon \geq 0$, the class of posterior beliefs inducing x as an ϵ -NE (WSNE) is a convex polytope described by $\text{poly}(n, m)$ linear inequalities.

Lemma 3.4. Given a family of non-empty polytopes $\mathcal{P}_1, \dots, \mathcal{P}_t \subseteq \Delta_M$ described by ℓ inequalities each, a point $\lambda \in \Delta_M$, and linear objectives $w_1, \dots, w_t \in \mathbb{R}^M$, the non-linear optimization problem (1) can be solved in $\text{poly}(t, \ell, M)$ time.

$$\begin{aligned} & \text{maximize} && \sum_{\sigma=1}^t \alpha_\sigma w_\sigma \cdot \mu_\sigma \\ & \text{subject to} && \alpha \in \Delta_t \\ & && \sum_{\sigma=1}^t \alpha_\sigma \mu_\sigma = \lambda \\ & && \mu_\sigma \in \mathcal{P}_\sigma, \quad \text{for } \sigma = 1, \dots, t. \end{aligned} \tag{1}$$

Before proving each of these lemmas, we first elaborate on how they imply Theorem 3.1. Given a signaling scheme φ with decomposition form (α, μ) , and an (approximate) equilibrium x^σ for each subgame corresponding to σ , the objective value is

$$f(\varphi, x) = \sum_{\sigma \in \Sigma} \alpha_\sigma \mathcal{F}(\mu_\sigma, x^\sigma)$$

where $\mathcal{F}(\mu, x)$ denotes $\mathbf{E}_{\theta \sim \mu} \mathbf{E}_{\vec{s} \sim x} \mathcal{F}(\theta, \vec{s})$.

Lemma 3.2 implies that, in order to complete the proof of Theorem 3.1, it suffices to show how to exactly optimize, in the claimed time, over signaling schemes in which $x^\sigma \in \mathcal{N}_{m, \epsilon}^n$ for each signal $\sigma \in \Sigma$. We may restrict attention to signaling scheme/equilibrium combinations for which each mixed strategy profile $x \in \mathcal{N}_{m, \epsilon}^n$ is selected for at most one subgame: when x is the equilibrium for both the subgames \mathcal{A}^{σ_1} and \mathcal{A}^{σ_2} , we can “merge” the two signals σ_1 and σ_2 into a signal (σ_1, σ_2) , giving rise to a new subgame $\mathcal{A}^{(\sigma_1, \sigma_2)}$ with posterior belief $\mu_{(\sigma_1, \sigma_2)} = \frac{\alpha_{\sigma_1}}{\alpha_{\sigma_1} + \alpha_{\sigma_2}} \mu_{\sigma_1} + \frac{\alpha_{\sigma_2}}{\alpha_{\sigma_1} + \alpha_{\sigma_2}} \mu_{\sigma_2}$ and

probability $\alpha_{(\sigma_1, \sigma_2)} = \alpha_{\sigma_1} + \alpha_{\sigma_2}$. Lemma 3.3 implies that x remains an (approximate) equilibrium of the merged subgame. Moreover, the objective is unchanged because $\mathcal{F}(\mu, x)$ is linear in its first argument.

For notational convenience we assume that each $x \in \mathcal{N}_{m, \epsilon}^n$ is induced as an equilibrium of exactly one subgame, by allowing signals which occur with probability 0, and discarding strategy profiles in $\mathcal{N}_{m, \epsilon}^n$ which can not be induced as equilibria of any posterior belief. The latter can be done in polynomial time, by checking whether the corresponding polytope (as given by Lemma 3.3) is empty. Writing $\mathcal{N}_{m, \epsilon}^n = \{x^1, \dots, x^t\}$ for $t = m^{O(\log m / \epsilon^2)}$, our optimization task can be written as follows.

$$\begin{aligned} & \text{maximize} && \sum_{\sigma=1}^t \alpha_{\sigma} \mathcal{F}(\mu_{\sigma}, x^{\sigma}) \\ & \text{subject to} && \alpha \in \Delta_t \\ & && \sum_{\sigma=1}^t \alpha_{\sigma} \mu_{\sigma} = \lambda \\ & && x^{\sigma} \text{ is an equilibrium of } \mathcal{A}^{\mu_{\sigma}}, \text{ for } \sigma = 1, \dots, t. \end{aligned} \tag{2}$$

Lemma 3.3, and the linearity of $\mathcal{F}(\cdot, \cdot)$ in its first argument, imply that optimization problem (2) is of the form given in (1) with $\ell = \text{poly}(n, m)$. Lemma 3.4, and our assumption that n is a constant, imply that (2) can be solved in time $\text{poly}(M, m^{\ln m / \epsilon^2})$. This completes the proof of Theorem 3.1.

3.1 Proof Sketch of Lemma 3.2

The special case of this lemma for ϵ -NE and $\delta = 0$ follows directly from the statement [LMM03, Theorem 2], by including an additional player in the game with no actions, and payoffs given by the objective tensor \mathcal{F} evaluated on the strategies of the other (real) players. The set $\mathcal{N}_{\epsilon, m}$ is taken to be the family of all mixed strategies which are uniformly distributed on a multiset contained in $[m]$ of size $\frac{3(n+1)^2 \ln(n+1)^2 m}{\epsilon^2}$.

More generally, at the heart of [LMM03, Theorem 2] is the fact that, for every mixed strategy profile $x \in \Delta_m^n$, one can choose a mixed strategy profile $y \in \mathcal{N}_{\epsilon, m}^n$, with $\text{supp}(y_i) \subseteq \text{supp}(x_i)$ for each $i \in [n]$, so that $|\mathcal{A}_i(j, y_{-i}) - \mathcal{A}_i(j, x_{-i})| \leq \frac{\epsilon}{2}$ for every player i and pure strategy $j \in [m]$. This has the implication that y is an $(\epsilon + \delta)$ -NE (WSNE) when x is a δ -NE (WSNE). Moreover, when the objective is included as a player in the game with no nontrivial strategies (which is without loss of generality), this also implies that $|\mathcal{F}(y) - \mathcal{F}(x)| \leq \epsilon$.

3.2 Proof of Lemma 3.3

For x to be an ϵ -NE of $\mathcal{A}^{\mu} = \sum_{\theta=1}^M \mu(\theta) \mathcal{A}^{\theta}$, for $\mu \in \Delta_M$, the following must hold:

$$\sum_{\theta=1}^M \mu(\theta) \mathcal{A}_i^{\theta}(x) \geq \sum_{\theta=1}^M \mu(\theta) \mathcal{A}_i^{\theta}(j, x_{-i}) - \epsilon, \quad \text{for } i \in [n], j \in [m]. \tag{3}$$

For an ϵ -WSNE, the analogous system of inequalities is:

$$\sum_{\theta=1}^M \mu(\theta) \mathcal{A}_i^{\theta}(j, x_{-i}) \geq \sum_{\theta=1}^M \mu(\theta) \mathcal{A}_i^{\theta}(k, x_{-i}) - \epsilon, \quad \text{for } i \in [n], j \in \text{supp}(x_i), k \in [m]. \tag{4}$$

Since x is fixed, in both cases we have a system $\text{poly}(n, m)$ linear inequalities in μ .

3.3 Proof of Lemma 3.4

We write an equivalent linear program via a change of variables. Specifically, we let $\gamma_{\sigma} = \alpha_{\sigma} \mu_{\sigma}$. Observe that after this change (1) becomes:

$$\begin{aligned}
& \text{maximize} && \sum_{\sigma=1}^t w_{\sigma} \cdot \gamma_{\sigma} \\
& \text{subject to} && \alpha \in \Delta_t \\
& && \sum_{\sigma=1}^t \gamma_{\sigma} = \lambda \\
& && \frac{\gamma_{\sigma}}{\alpha_{\sigma}} \in \mathcal{P}_{\sigma}, \quad \text{for } \sigma = 1, \dots, t.
\end{aligned} \tag{5}$$

(5) is not yet a linear program. However, note that the constraint $\alpha \in \Delta_t$ is implied if we simply add the constraints $\gamma_{\sigma} \geq 0$. Moreover, because $\gamma_{\sigma}/\alpha_{\sigma} \in \Delta_M$, $\alpha_{\sigma} = \sum_{\theta} \gamma_{\sigma}(\theta)$ holds for every feasible solution, allowing us to simplify the constraint $\gamma_{\sigma}/\alpha_{\sigma} \in \mathcal{P}_{\sigma}$. Since \mathcal{P}_{σ} is described an explicit linear system $A^{\sigma}y \preceq b^{\sigma}$, the non-linear system of inequalities $A^{\sigma}\gamma_{\sigma}/\alpha_{\sigma} \preceq b^{\sigma}$ can be re-written as the linear system $A^{\sigma}\gamma_{\sigma} \preceq (\sum_{\theta} \gamma_{\sigma}(\theta))b^{\sigma}$. This results in an equivalent linear program with variables $\gamma_1, \dots, \gamma_t \in \mathbb{R}_+^M$, from which $\alpha_{\sigma} = \sum_{\theta} \gamma_{\sigma}(\theta)$ and $\mu_{\sigma} = \gamma_{\sigma}/(\sum_{\theta} \gamma_{\sigma}(\theta))$ can be recovered efficiently.

3.4 Remarks

Zero-sum games When applied to two-player zero-sum games with the objective to maximize the first-player's payoff, our signaling scheme provides a stronger guarantee. In such setting, both players retain the same payoff in any exact Nash equilibrium. Also, any ϵ -equilibria gives a payoff that is ϵ close to the payoff of any exact equilibrium. Thus, the signaling scheme provided in Theorem 3.1 can be directly compare with the quality of the optimal signaling scheme without concerning equilibrium selection, instead of a bi-criteria guarantee.

Reducing the number of signals Although the signaling scheme provided in Theorem 3.1 might use quasi-polynomial number of signals, we can reduce the number of signals to $M + 1$. Let w_{λ} be the objective value of the signaling scheme, and consider the set of t signals used $\mu_1 \cdots \mu_t$ and their corresponding expected objective values $w_1 \cdots w_t$. Observe that the $M + 1$ dimension point (w_{λ}, λ) is a convex combination of the set of points $P = \{(w_1, \mu_1) \cdots (w_t, \mu_t)\}$. Since w_{λ} is maximized, (w_{λ}, λ) belongs to some facet of the convex hull of P . Hence by Carathéodory's theorem, (w_{λ}, λ) can be written as a convex combination of only $M + 1$ points from P , and such decomposition can be computed in time polynomial in the size of P . This decomposition gives a valid signaling scheme with the same objective value, using only $M + 1$ signals.

Stackelberg games Our result can be extended to *Stackelberg games* which often arise in security games [Dug14]. Recall that in a Stackelberg game [vSBHU10], one player (the leader) first commits to a (mixed) strategy, and then all other players (followers) simultaneously play their strategies upon learning the leader's strategy. Our result can be readily extended to Bayesian Stackelberg games when the objective of the signaling scheme is to maximize leader's payoff. In this case, we can simply drop the constraints regarding the leader in the polytopes defined in Lemma 3.3, and only require the followers to play an approximate equilibrium in our algorithm presented in Theorem 3.1.

4 Signaling in Bayesian Single Item Auctions

In this section, we consider the signaling problem in a Bayesian single item auction, as described in Section 2. As in [DIR14], our goal is to maximize the social welfare subject to a communication constraint k on signaling scheme. We present a quasi-polynomial time approximation scheme for this problem when the valuation distribution \mathcal{D} is given explicitly, and as a corollary also when \mathcal{D} is given by a sampling oracle. To aid in our proof, we begin with some technical background drawn

from related work. In the discussion of this section, we will use (n, M, k, r) -second-price signaling to denote the welfare-maximization signaling problem with n bidders, M configurations of the item, k signals, and an explicitly-described valuation distribution supported on r matrices.

4.1 Background: Reducing to Submodular Maximization for small r

Recall that a signaling scheme is a randomized map from states of nature to signals. A deterministic signaling scheme with k signals, naturally, is a deterministic map $\varphi : \Theta \rightarrow [k]$. Equivalently, such a scheme corresponds to a partition of the states of nature Θ into k classes $\Theta_1, \dots, \Theta_k$, one per signal. As might be apparent from Section 3, in general games randomized signaling schemes may outperform their deterministic counterparts. However, this is not the case in our auction setting.

Lemma 4.1 ([DIR14]). *For communication-constrained signaling in Bayesian second-price auctions, there always exists a deterministic signaling scheme which maximizes expected social welfare.*

Thus instead of solving the signaling problem by assigning probabilities in a continuous domain, we can exploit the fact that our signaling problem has a combinatorial solution. This is the basis for the $(1 - 1/e)$ -approximation algorithm of [DIR14] for (n, M, r, k) -second price signaling, which we outline next. This algorithm reduces the signaling problem to submodular optimization,⁵ and is computationally efficient when r , the support size of the valuation distribution \mathcal{D} , is small. Using $\mathcal{V}^1, \dots, \mathcal{V}^r$ and ρ_1, \dots, ρ_r to denote the support of \mathcal{D} and the corresponding probabilities, respectively, consider a deterministic signaling scheme in partition form $\Theta_1, \dots, \Theta_k$. Such a signaling scheme induces kr subgames, one for each pair $(\Theta_\sigma, \mathcal{V}^t)$. In each such subgame, there is a unique winning player of the auction — the player i with maximum posterior value

$$\frac{1}{\lambda(\Theta_\sigma)} \sum_{\theta \in \Theta_\sigma} \lambda(\theta) \mathcal{V}^t(i, \theta).$$

The social welfare, therefore, is given by

$$\sum_{t=1}^r \sum_{\sigma=1}^k \max_{i \in [n]} \sum_{\theta \in \Theta_\sigma} \rho_t \lambda(\theta) \mathcal{V}^t(i, \theta). \quad (6)$$

The key observation behind the algorithm of [DIR14] is that, given a “guess” for the winner of the auction in each of these kr subgames, the optimal signaling scheme $(\Theta_1, \dots, \Theta_k)$ can be recovered in $\text{poly}(n, M, r)$ time. More generally, given a list w_1, \dots, w_k , where $w_\sigma \in [n]^r$ is the *winner tuple* for signal σ , we can write the expected social welfare of a deterministic signaling scheme $\Theta_1, \dots, \Theta_k$ under the assumption that $w_\sigma(t)$ wins in the event (σ, t) .

$$\text{welfare}(\Theta_1, \dots, \Theta_k, w_1, \dots, w_k) = \sum_{t=1}^r \sum_{\sigma=1}^k \sum_{\theta \in \Theta_\sigma} \rho_t \lambda(\theta) \mathcal{V}^t(w_\sigma(t), \theta). \quad (7)$$

It is now clear that the optimal scheme, assuming winner tuples w_1, \dots, w_k , maps θ to the signal σ maximizing $\sum_{t=1}^r \sum_{\sigma=1}^k \rho_t \lambda(\theta) \mathcal{V}^t(w_\sigma(t), \theta)$. The resulting scheme has social welfare at least as given by the following expression, with equality holding for the optimal set of guesses w .

$$\text{welfare}(w_1, \dots, w_k) = \sum_{\theta \in \Theta} \max_{\sigma \in [k]} \sum_{t=1}^r \rho_t \lambda(\theta) \mathcal{V}^t(w_\sigma(t), \theta). \quad (8)$$

We summarize this discussion by the following lemma.

⁵Recall that a function $f : 2^\Omega \rightarrow \mathbb{R}$ on a ground set Ω is *submodular* if it satisfies $f(X) + f(Y) \geq f(X \cup Y) + f(X \cap Y)$ for every $X, Y \subseteq \Omega$.

Lemma 4.2 ([DIR14]). *Consider an instance (λ, \mathcal{D}) of (n, M, r, k) -second-price signaling. Given a list $w_1, \dots, w_k \in [n]^r$ of winner tuples, a deterministic signaling scheme $\varphi : \Theta \rightarrow [k]$ with $\text{welfare}(\varphi, \mathcal{D}, \lambda) \geq \text{welfare}(w_1, \dots, w_k)$ can be computed in $\text{poly}(n, M, r)$ time. Moreover, there is a set of winner tuples w_1^*, \dots, w_k^* with $\text{welfare}(w_1^*, \dots, w_k^*)$ equal to the optimum welfare of a k -signal scheme.*

The above lemma naturally leads to a set function maximization problem subject to a cardinality constraint of k . By symmetry, we can think of a list of winner tuples w_1, \dots, w_k equivalently as an (unordered) set $W = \{w_1, \dots, w_k\}$, assigning an order arbitrarily. Assuming winning tuples W , the signaling scheme of Lemma 4.2 has welfare at least

$$\text{welfare}(W) = \sum_{\theta \in \Theta} \max_{w \in W} \sum_{t=1}^r \rho_t \lambda(\theta) v^t(w(t), \theta), \quad (9)$$

with equality holding at optimality. The above function naturally extends to all $W \subseteq [n]^r$, and is monotone nondecreasing and submodular. To see this, note that it can be written in the form $\sum_{\theta} \max_{w \in W} g_{\theta}(w)$ — the function $\max_{w \in W} g_{\theta}(w)$ is monotone and submodular for each fixed θ , and therefore the sum of such functions is as well. Thus maximizing the social welfare subject to communication constraint k can be formulated as the following submodular maximization problem subject to a cardinality constraint:

$$\max_{|W|=k, W \subseteq [n]^r} \text{welfare}(W). \quad (10)$$

The problem of maximizing a monotone submodular function subject to cardinality constraint is NP-hard in general, though can be approximated to within a factor of $(1 - 1/e)$ using a simple greedy algorithm [NWF78]. The runtime of the algorithm is polynomial in the size of the ground set, as well as the time needed to evaluate the function. This leads to a deterministic $(1 - 1/e)$ -approximation algorithm for (n, M, k, r) -second-price auction signaling which runs in time $\text{poly}(n^r, M)$.

Theorem 4.3 ([DIR14]). *There is a deterministic $\text{poly}(n^r, M)$ -time $(1 - \frac{1}{e})$ -approximation algorithm for welfare maximization in (n, M, k, r) -second-price auction signaling.*

Finally, we note that the above approximation guarantee is optimal, assuming $P \neq NP$, even when $r = 1$ [DIR14]. Thus, our quasi-polynomial-time algorithms necessarily must yield no better than a $(1 - 1/e)$ -approximation, unless NP-complete problems can be solved in quasi-polynomial time.

4.2 Quasi-Polynomial Algorithm for Auction Signaling

The exponential dependence on r in Theorem 4.3 limits its applicability. In this section, we present an approximation algorithm for the (M, n, r, k) -signaling problem with runtime exponential in $\log(nr)$ instead — quasi-polynomial in the size of the explicit input representation. The algorithm achieves the same multiplicative approximation guarantee of $(1 - 1/e)$, modulo an additive loss that can be made arbitrarily small. As a consequence, we get a randomized quasi-polynomial-time algorithm with a similar guarantee when valuation distributions are given by a sampling oracle.

Theorem 4.4. *Given an instance of (n, M, k, r) -second-price signaling in the explicit model, and an approximation parameter $\epsilon > 0$, there is a deterministic algorithm which runs in $M^{O(\log(nr)/\epsilon^2)}$ -time, and computes a signaling scheme with k signals and expected social welfare $(1 - 1/e)(\text{OPT} - \epsilon)$, where OPT denotes the optimal welfare of a k -signal scheme.*

Corollary 4.5. *Given an instance of (n, M, k, r) -second-price signaling where the valuation distribution is given via a sampling oracle, and parameters $\epsilon, \delta > 0$, there is a Monte Carlo algorithm which runs in time $M^{O(\log(nM\delta^{-1}\epsilon^{-2})/\epsilon^2)}$, and with success probability $1 - \delta$ outputs a signaling scheme with k signals and expected social welfare $(1 - 1/e)(OPT - 2\epsilon)$.*

Corollary 4.5 follows from Theorem 4.4 by using standard Monte Carlo sampling and convergence arguments, exploiting the fact that there are at most k^M deterministic signaling schemes. We include a self contained proof in Section 4.4. Theorem 4.4, on the other hand, follows from the following Lemma, which we will prove in Section 4.3.

Lemma 4.6. *Given an instance of the (n, M, k, r) -second-price signaling problem, and a parameter $\epsilon > 0$, there is a family of winner tuples $\mathcal{N}_\epsilon \subseteq [n]^r$ with $|\mathcal{N}_\epsilon| = M^{O(\log(nr)/\epsilon^2)}$, satisfying*

$$\max_{|W|=k, W \subseteq \mathcal{N}_\epsilon} \text{welfare}(W) \geq \max_{|W|=k, W \subseteq [n]^r} \text{welfare}(W) - \epsilon. \quad (11)$$

Moreover, \mathcal{N}_ϵ can be enumerated in time $M^{O(\log(nr)/\epsilon^2)}$.

Armed with Lemma 4.6, we complete the proof of Theorem 4.4. We follow the blueprint outlined in Section 4.1, using \mathcal{N}_ϵ as our ground set of winner tuples in lieu of $[n]^r$. Specifically, we use the greedy algorithm to approximately maximize $\text{welfare}(W)$ over $W \subseteq \mathcal{N}_\epsilon$ with $|W| = k$. This leads to a signaling scheme with welfare at least $(1 - 1/e)(OPT - \epsilon)$. Given the reduced size of the ground set of our submodular function, the runtime is now $\text{poly}(|\mathcal{N}_\epsilon|, n, M, r) = M^{O(\log(nr)/\epsilon^2)}$.

4.3 Proof of Lemma 4.6

Given a multiset of item configurations $Y \subseteq \Theta$, we distinguish the winner tuple *induced* by Y , which simply lists the winning player for each $V^t \in \text{supp}(\mathcal{D})$ assuming θ is drawn uniformly from Y . Formally, the winner tuple $w(Y) = (w_1(Y), \dots, w_r(Y))$ induced by Y is defined by $w_t(Y) = \arg\max_{i=1}^n \mathbf{E}_{\theta \sim Y}[\mathcal{V}^t(i, \theta)]$. We then define our set \mathcal{N}_ϵ as follows.

$$\mathcal{N}_\epsilon = \left\{ w(Y) : |Y| = \frac{2 \ln(4nr)}{\epsilon^2}, Y \subseteq \Theta \right\}$$

\mathcal{N}_ϵ is of the claimed size. It remains to show that considering only winner tuples in \mathcal{N}_ϵ in the optimization problem $\max_{|W|=k} \text{welfare}(W)$ results in a loss of welfare no more than an additive ϵ . This hinges on the following claim.

Claim 4.7. *Fix $\epsilon > 0$. For each $X \subseteq \Theta$, there is multiset $Y \subseteq X$ with $|Y| = 2 \ln(4nr)/\epsilon^2$ such that,*

$$\left| \mathbf{E}_{\theta \sim \lambda} [\mathcal{V}^t(i, \theta) | \theta \in X] - \mathbf{E}_{\theta \sim Y} [\mathcal{V}^t(i, \theta)] \right| \leq \epsilon/2$$

for every $t = 1, \dots, r$ and $i = 1, \dots, n$.

Proof. Consider taking $2 \ln(4nr)/\epsilon^2$ samples Y from the distribution $\lambda|X$. Standard application of the Hoeffding bound and the union bound shows that Y satisfies the desired guarantee with probability at least $1/2 > 0$. The existential result then follows. \square

As a consequence, for every deterministic signaling scheme $(\Theta_1, \dots, \Theta_k)$, we can choose a tuple of multisets (Y_1, \dots, Y_k) with $Y_\sigma \subseteq \Theta_\sigma$ and $|Y_\sigma| = 2 \ln(4nr)/\epsilon^2$ for each $\sigma \in [k]$, so that the empirical distribution Y_σ always approximates, to within an additive $\epsilon/2$, the posterior expected value of every player in subgame σ . Hence, using the empirical distribution Y_σ to determine the winner of each signal Θ_σ preserves the social welfare in the corresponding subgame up to an additive ϵ . When $(\Theta_1, \dots, \Theta_k)$ is an optimal scheme, this implies that $\text{welfare}(w(Y_1), \dots, w(Y_k)) \geq OPT - \epsilon$, as needed.

4.4 Proof of Corollary 4.5

We fix an arbitrary deterministic signaling scheme φ (i.e. partition on items), and consider the randomness in valuation matrices. Let $f(\varphi, V)$ be the expected welfare induced by φ when the valuation matrix is V , and let $f(\varphi, \mathcal{D})$ be the expected welfare induced by φ when the valuation is drawn from \mathcal{D} . By definition, $f(\varphi, \mathcal{D}) = E_{V \sim \mathcal{D}}[f(\varphi, V)] \in [0, 1]$. In order to estimate $f(\varphi, \mathcal{D})$, we take r samples V_1, \dots, V_r from \mathcal{D} and apply our algorithm in Theorem 4.4 as if the valuation distribution is uniform on V_1, \dots, V_r . We apply Hoeffding’s inequality to bound the probability that welfare estimated from V_1, \dots, V_r is far from its expectation.

$$Pr \left(\left| \frac{1}{r} \sum_{j=1}^r f(\varphi, V_j) - f(\varphi, \mathcal{D}) \right| \geq \epsilon/2 \right) \leq 2e^{-r\epsilon^2/2},$$

so by setting $r = \Theta((M \log k + \log(\delta^{-1}))/\epsilon^2)$, we get this probability to be at most δk^{-M} . Because there are at most k^M deterministic signaling schemes, by union bound, we preserve the welfare up to an additive error $\epsilon/2$ for all of them. We conclude that, in this situation, our signaling scheme computed as in Theorem 4.4 has expected welfare at least $(1 - 1/e)(OPT - 2\epsilon)$.

References

- [Ake70] George A Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The quarterly journal of economics*, pages 488–500, 1970.
- [BGSZ03] Bruno Basset, Olivier Gossner, Marco Scarsini, and Shmuel Zamir. Positive value of information in games. *International Journal of Game Theory*, 32(1):17–31, 2003.
- [Bla51] David Blackwell. Comparison of experiments. In *Second Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 93–102, 1951.
- [BMS12] Peter Bro Miltersen and Or Sheffet. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 234–247, 2012.
- [DIR14] Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. Constrained signaling in auction design. In *Proceedings of the 25th ACM Symposium on Discrete Algorithms (SODA)*, 2014.
- [Dug14] Shaddin Dughmi. On the hardness of signaling. In *Proceedings of the 55th IEEE Symposium on Foundations of Computer Science (FOCS)*, 2014. To appear.
- [EFG⁺12] Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. Signaling schemes for revenue maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 514–531, 2012.
- [GD13] Mingyu Guo and Argyrios Deligkas. Revenue maximization via hiding item attributes. In *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, pages 157–163. AAAI Press, 2013.
- [Hir71] Jack Hirshleifer. The private and social value of information and the reward to inventive activity. *The American Economic Review*, 61(4):561–574, 1971.

- [KG09] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. Technical report, National Bureau of Economic Research, 2009.
- [LMM03] Richard Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic commerce (EC)*, pages 36–41, 2003.
- [LRS10] Ehud Lehrer, Dinah Rosenberg, and Eran Shmaya. Signaling and mediation in games with common interests. *Games and Economic Behavior*, 68(2):670–682, 2010.
- [MW82] Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5), 1982.
- [NWF78] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher. An analysis of approximations for maximizing submodular set functions – I. *Mathematical Programming*, 14(3):265–294, 1978.
- [SKT13] Vasilis Syrgkanis, David Kempe, and Eva Tardos. Information asymmetries in common-value auctions with discrete signals. Preprint., 2013.
- [Spe73] Michael Spence. Job market signaling. *The quarterly journal of Economics*, 87(3):355–374, 1973.
- [vSBHU10] Heinrich von Stackelberg, Damien Bazin, Rowland Hill, and Lynn Urch. *Market structure and equilibrium*. Springer, 2010.